

BASE-RATE NEGLECT AS A FUNCTION OF BASE RATES IN PROBABILISTIC CONTINGENCY LEARNING

FLORIAN KUTZNER, PETER FREYTAG, TOBIAS VOGEL, AND KLAUS FIEDLER

UNIVERSITY OF HEIDELBERG

When humans predict criterion events based on probabilistic predictors, they often lend excessive weight to the predictor and insufficient weight to the base rate of the criterion event. In an operant analysis, using a matching-to-sample paradigm, Goodie and Fantino (1996) showed that humans exhibit base-rate neglect when predictors are associated with criterion events through physical similarity. In partial replications of their studies, we demonstrated similar effects when the predictors resembled the criterion events in terms of similarly skewed base rates. Participants' predictions were biased toward the more (or less) frequent criterion event following the more (or less) frequent predictor. This finding adds to the growing evidence for pseudocontingencies (Fiedler & Freytag, 2004), a framework that stresses base-rate influences on contingency learning.

Key words: pseudocontingency, skewed base rates, base-rate fallacy, probabilistic contingency learning, matching-to-sample, humans, computer keyboard

When predicting criterion events from predictors in probabilistic settings, it is normatively appropriate to consider two kinds of information, the global base rate of the criterion events and the case-specific predictor values. However, whereas predictor information is readily utilized, even when the predictor's actual validity is questionable, the information inherent in the criterion-event base rate is often not exploited sufficiently (Tversky & Kahneman, 1982). For example, when the predictor is an eyewitness's testimony that a suspect's car was blue, human participants tend to believe that it actually was blue, even in the face of evidence that, in that particular locale, the base rate of blue cars is low. In this example, they underweigh the base rate of the criterion events (the ratio of colors in that locale) and act as if only the predictor (the testified-to color) provided relevant information.

In an intriguing line of research, Goodie and Fantino (1996, 1999) translated the base-rate neglect phenomenon into an operant learning paradigm. On every trial, a predictor

(e.g., either a high or a low tone) is presented and participants have to choose between two possible criterion events, for example, pressing one of two keys on a keyboard. The choice is reinforced if pressing the chosen key was the criterion event and it is punished if pressing the other key was the criterion event. Only choice together with reinforcement indicates to the participant which key press actually was the criterion event. In this paradigm, the relevant base-rate information lies in the reinforcement rate associated with pressing the two keys. Case-specific information conveyed by the predictor depends on the contingency between it and the criterion event. That is, it depends on the reinforcement rate associated with one predictor minus the reinforcement rate associated with the other predictor (Allan, 1993).

Goodie and Fantino (1999) found that both base-rate and predictor information influenced choice, but participants were more sensitive to the latter than to the former. Additionally, even when the predictors did not contain any useful statistical information, because of lack of contingency with the criterion events, there was a bias toward choosing the criterion events that superficially resembled the predictors (Goodie & Fantino, 1996). For instance, when blue and green rectangles were used as predictors and blue and green rectangles were associated with the response options, there was a tendency to select the response options by matching their

The research underlying this paper was supported by a grant from the Deutsche Forschungsgemeinschaft (DFG) awarded to the second and last authors.

Correspondence should be addressed to Florian Kutzner, Department of Psychology, University of Heidelberg, Hauptstrasse 47-51, 69117 Heidelberg, Germany, Tel +49-6221-547366, Fax +49-6221-547745 (e-mail: florian.kutzner@psychologie.uni-heidelberg.de).

doi: 10.1901/jeab.2008.90-23

colors to the colors named by the predictors. Moreover, this tendency distracted participants from choosing the option with the higher base rate of reinforcement.

Thus, in what we interpret as base-rate neglect, different kinds of predictor information cause deviations from Bayesian standards for the use of criterion-event base rates. Choice may overstate statistical information when a real contingency exists between the predictor and the criterion events. Or, alternatively, even in the absence of useful statistical information, choice may reflect “intensional” information, such as the physical similarity of predictors and criterion events. Adding to these influences, we show that the similarity of predictor and criterion-event base rates prevents the normative impact of the criterion-event base rate. We suggest that choice in an operant paradigm is influenced by what we refer to as pseudocontingency (PC; see Fiedler & Freytag, 2004).

A PC is a contingency judgment based only on statistically inappropriate base-rate information—hence the term “pseudocontingency.” It reverses the finding that contingencies override base rates. A growing body of evidence shows that when the base rates of two uncorrelated variables are jointly skewed, they will be judged as correlated (see Fiedler, Freytag, & Unkelbach, 2007). This finding has implications for a prediction task when one predictor is more frequent than the other, and one criterion event is also more frequent than the other. In this situation the more frequent criterion event is more likely to be predicted from the more frequent predictor, whereas the less frequent criterion event will be more likely predicted from the less frequent predictor. In other words, the similarity in base rates is taken for a contingency.

For example, imagine you just moved to a new town, where you are confronted with choosing between two highways that lead to your office and that take turns in being congested. Every day, your task is to predict which is the light-traffic highway. By and by you learn that one of the highways has the higher rate of light traffic, so you tend to take that highway. However, you might still take case-specific predictors into account. For example, you might take into account whether you are traveling during the week, as you usually do, or on a weekend, which happens

less often. This example leaves two base rates skewed by your experience. One predictor, “weekday”, is more frequent than the other, “weekend,” and one of the highways has the higher base rate of the criterion event, “light-traffic highway.” Given these two base rates, a PC may exist between the day of the week and which is the light-traffic highway. If it did, you are more likely to choose the frequently light-traffic highway (the more frequent criterion event) on a weekday (the more frequent predictor) than on a weekend (the less frequent predictor). Without information about the actual contingency between the predictor and the criterion event, a PC links the higher base-rate events (predictor and criterion event) to each other and the lower base-rate events to each other. Thus, a PC can be understood as a statistically unwarranted analogical inference: What is frequent (or infrequent) in one dimension should also be frequent (or infrequent) in another dimension.

In our attempt to apply the pseudocontingency framework to an operant paradigm, we used a matching-to-sample (MTS) task similar to the one of Goodie and Fantino (1996, 1999). Participants were repeatedly presented with one of two predictors, a high tone and a low tone, in response to which they were asked to choose between two criterion events by pressing one of two keys on a computer keyboard. Every trial produced feedback as to whether the choice was correct or incorrect. Monetary rewards and punishers of equal size accompanied the feedback. Thus, every trial involved the participant’s prediction about which of the two criterion events (key presses) was present; and choosing the correct criterion event resulted in reinforcement whereas choosing the incorrect criterion event resulted in punishment. Importantly, the actual contingency between predictors and criterion events was set to zero.

In three experiments, we created conditions where the base rates of both the predictors and the criterion events were skewed. We expected the participants’ behaviors to be a function of two tendencies working together. One tendency is to choose the criterion event with the higher base rate of occurrence. Previous MTS studies provided evidence for probability matching rather than optimizing. That is, choice proportions for two options

roughly equaled the reinforcement rates associated with the options (see Humphreys, 1939; Shanks, 1990), rather than showing an exclusive preference for the more frequently reinforced option. With monetary incentives, a tendency toward optimizing (i.e., exclusively choosing the more frequently reinforced option) has also been reported (Shanks, Tunney, & McCarthy, 2002). In view of these mixed findings, we expected the choice proportions to be between probability matching and optimizing.

We also expected a PC to emerge, that is, a bias toward the criterion event with a base rate of occurrence similar to the base rate of the predictor. Specifically, this implies a higher choice proportion for the more frequent criterion event on trials involving the more frequent predictor.

EXPERIMENT 1

We used two computer-generated tones (a high-pitch piano tone and low-pitch saxophone tone) as predictors and two keys ("A" and "Ä") on the left and on the right side of a German computer keyboard, respectively, as response options. Thus a preexisting association in terms of physical similarity between tones and response options was extremely unlikely.

METHOD

Participants

Twenty-four undergraduate students (20 females, 4 males) from the University of Heidelberg participated. They signed up on a placard promoting an information-processing experiment with an ordinary performance-contingent reward.

Apparatus

Up to 6 students participated at the same time in an artificially lit laboratory located within the psychology department. They were seated in cubicles equipped with a personal computer without being able to observe the other participants. A commercial software program (KB Piano 1.2[®]) was used to generate the tones, which were delivered via earphones. Participants could adjust the volume to their liking.

Procedure

Participants were randomly assigned to one of the two groups with differing base rates and every participant was exposed to the two different tones. This resulted in a 2 (base rates skewed vs. no skew) \times 2 (more frequent tone vs. less frequent tone) mixed design with repeated measures on the latter factor. The computer delivered the instructions, controlled stimulus presentation, and recorded participants' choices.

Participants were instructed to find out which key was correct following the presentation of each of the tones. Trials started with the presentation of a tone that lasted for a maximum of 2.5 s. The keyboard was locked for 1 s after the onset of the tone. Once the keyboard was unlocked, participants could terminate the tone by pressing one of the keys. This produced immediate feedback by displaying the German words for correct and incorrect for 1.5 s. Thereafter the next tone was presented.

Sessions lasted until participants had responded to a random sequence of 160 tones. On average, sessions lasted for about 13 min. Only one session was conducted per participant. Participants began each session with an account of 3€ (Euros; approximately \$4 US) of prospective compensation. For each correct choice 0.05€ was added to this account. For each incorrect choice 0.05€ was subtracted. At the end of each trial, the outcome on the current trial (either plus or minus 0.05€), the choice (either left or right), and the updated account value were displayed on the screen for 1.5 s together with the feedback. Participants were compensated following the session.

For every participant, the computer generated a random sequence of tones and corresponding correct choices by drawing without replacement from one of two predetermined distributions. In the skewed distribution (top panel in Table 1) the high-pitch tone was three times as frequent as the low-pitch tone, and the left key was reinforced three times as frequently as the right key. The distribution of tones and keys was counterbalanced. That is, for an equal number of participants the left or the right key was more frequently reinforced, and the high or the low-pitch tone was more frequent. In the no-skew distribution (middle panel in Table 1) the tones and correct choices were evenly distributed. The actual

Table 1

Contingency tables indicating the predetermined distributions used in the experiments.

		Reinforced option		
		Left key	Right key	
Experiment 1 & 2				
Skewed condition				
Predictor	High-pitch tone	90	30	120
	Low-pitch tone	30	10	40
		120	40	160
No-skew condition				
Predictor	High-pitch tone	40	40	80
	Low-pitch tone	40	40	80
		80	80	160
Experiment 2				
Criterion-skewed condition				
Predictor	High-pitch tone	60	20	80
	Low-pitch tone	60	20	80
		120	40	160

Note. In Experiment 1, the base rates of both predictors (tones) and criterion events (reinforced options) were either skewed at a ratio of 3:1 (top panel) or evenly distributed at a ratio of 1:1 (middle panel). Experiment 2 included an additional condition in which the criterion-event base rates were skewed at a ratio of 3:1, whereas the predictor base rates were evenly distributed (1:1; bottom panel).

contingency between the tones and the reinforced options was always zero, that is, the proportions of reinforcement after each tone did not differ as a function of the tone that had been presented.

Note that the implications of the reinforcement distributions remain unchanged when punishment is considered. Swapping the columns of the matrices in Table 1 results in the distributions for punished options. The resulting distributions share the same properties in terms of actual contingencies. However, in the skewed condition, the other response option is associated with the higher frequency of punishment. This implies a PC in the opposite direction. That is, the proportion of punishment for the more frequently punished choice should be judged higher, after the more frequent tone. Given that participants try to maximize reinforcement, their choices should then follow the inverse contingency, which is identical to the contingency when reinforcers are considered. Thus, our hypothesis remains unchanged whether focused on reinforcement or punishment.

RESULTS AND DISCUSSION

For every participant in both groups, the mean conditional choice proportions for the more frequently reinforced option, given the more frequent and less frequent tones, were

calculated. We analyzed only the second half of the trials, thus excluding the variability that occurred during early trials¹. Table 2 shows the difference between these choice proportions. Positive values represent a bias toward the more frequently reinforced option following presentation of the more frequent tone compared to the less frequent tone. In the skewed condition, 11 of 12 participants showed a positive bias. In the no-skew condition, positive and negative biases occurred with similar frequencies (5 versus 7 participants, respectively).

The conditional choice proportions were additionally submitted to a two-factor repeated-measures analysis of variance with distribution type as the between-subject factor (skewed vs. no skew) and tone type as the within-subject factor (more frequent vs. less frequent). The analysis revealed a large main effect for type of distribution, $F(1, 22) = 38.35, p < .01$, a main effect for type of tone, $F(1, 22) = 4.93, p < .05$, and an interaction, $F(1, 22) = 4.57, p < .05$. Figure 1 shows the mean choice proportions

¹ Additional analyses for the second half of the trials of the skewed condition revealed close to equal proportions of reinforcement, $t(11) = 1.12, p = .27$, after the more frequently reinforced option ($M = 3.31, SD = 1.14$) and after the less frequently reinforced option ($M = 2.95, SD = 1.28$) and no correlation with the differences in participants' choice proportions, $r(12) = .06, p = .80$.

Table 2

Bias towards choosing the more frequently reinforced option at a higher rate after the more frequent tone than after the less frequent tone in Experiment 1.

Conditions	Skewed		No skew	
Individual Data	-.15	(.72, .87)	-.15	(.43, .58)
	.02	(.97, .94)	-.13	(.40, .53)
	.05	(.90, .85)	-.08	(.63, .71)
	.06	(.98, .92)	-.07	(.50, .57)
	.07	(.89, .82)	-.05	(.48, .53)
	.07	(.60, .52)	-.04	(.56, .59)
	.09	(.93, .84)	-.03	(.58, .60)
	.14	(.91, .77)	.01	(.53, .52)
	.15	(.95, .80)	.08	(.59, .51)
	.21	(.79, .58)	.09	(.66, .56)
	.29	(.86, .57)	.13	(.70, .57)
	.53	(.92, .39)	.26	(.59, .33)
Mean	.11		-.02	
SD	.15		.14	

Note. The data for each participant are ordered from lowest to highest value of bias. The data are means from the second half of the trials. The numbers in parentheses are the conditional proportions of choosing the more frequently reinforced option after the presentation of the more frequent tone and after the less frequent tone. Bias is defined as the difference between the choice proportion for the more frequently reinforced option after the more frequent tone minus the choice proportion for the more frequently reinforced option after the less frequent tone.

for the more frequently reinforced option, conditional on the type of tone for both groups.

The main effect for distribution type shows that participants were sensitive to the different base rates of reinforcement for the two options. When averaged across tone types, participants in the skewed condition chose the more frequently reinforced option at a slightly higher rate than its 75% reinforcement rate ($M = .80$, $SD = 0.12$). In the no-skew condition their preference was slight ($M = .55$, $SD = .07$). However (and crucially), in the skewed condition the degree to which choice was governed by the reinforcement base rate depended on the tone, $t(11) = 2.68$, $p < .05$. The choice proportion of the more frequently reinforced option was higher after the more frequent tone ($M = .87$, $SD = .11$) than after the less frequent one ($M = .74$, $SD = .18$).

The results were consistent with the PC phenomenon. A majority of participants behaved as though there was a contingency between the predictor and the criterion event consistent with the skew of the base rates. Moreover, they did so in the absence of an

actual contingency and of physical similarity between predictors and criterion events.

However, because we jointly manipulated the predictor and criterion-event base rates, we could not determine how the two base-rate distributions contributed to participants' choice behavior. Therefore, in Experiment 2, we included a condition in which only the criterion-event base rate was skewed.

EXPERIMENT 2

Experiment 2 was an extended replication of Experiment 1 with an additional base-rate distribution and slightly modified instructions. This time instructions focused on maximizing the returns instead of finding which was the correct key given the two tones. In addition to the distributions used in Experiment 1, we included a new distribution (see the bottom panel of Table 1) in which the reinforcement base rate for the response options, but not the tone base rate, was skewed at the ratio of three to one. In this condition we hypothesized that participants would choose the more frequently reinforced option at the same rate following either tone. By contrast, when the tone base rate was also skewed, we expected a higher choice proportion for the more frequently reinforced option on trials involving the more frequent tone than on trials involving the less frequent one.

METHOD

Participants and Design

Forty-two students (12 female, 30 male) from the University of Mannheim participated in the study. They were part of an experimental pool and signed up to participate in an experiment with an ordinary, performance-contingent reward without having been informed about the content of the study.

Apparatus

Up to 15 students participated at the same time in a naturally lit laboratory located within the psychology department. They were seated next to each other in front of personal computers, unable to observe the other participants' choices. The tones were delivered via earphones, and participants could adjust the volume to their liking.

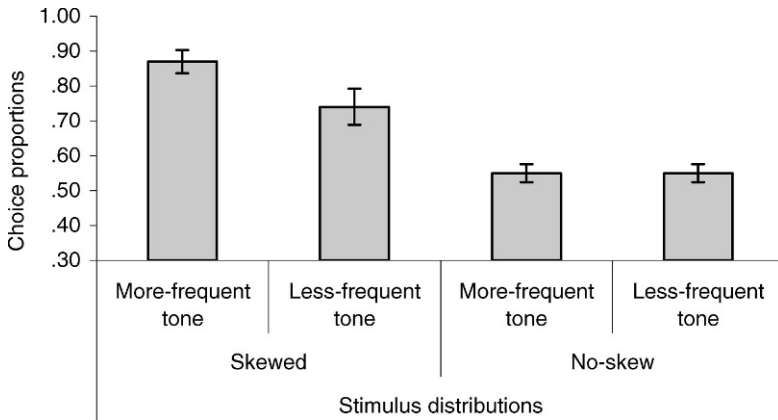


Fig. 1. Mean conditional choice proportions for the more frequently reinforced option as a function of the type of tone (more frequent or less frequent) on the trial. Error bars represent the standard error of the mean.

Procedure

Participants were randomly assigned to one of three groups, and every participant was exposed to the two different tones, resulting in a 3 (skewed vs. criterion-skewed vs. no-skew) \times 2 (more frequent tone vs. less frequent tone) mixed design with repeated measures on the latter factor. The computer delivered the instructions, controlled stimulus presentation, and recorded participants' choices. Participants were instructed to maximize their returns. The task was identical to that of Experiment 1. Only one session was conducted per participant.

RESULTS AND DISCUSSION

Again, choice proportions for the more frequently reinforced option conditional on the type of tone in the trial were calculated for the second half of the trials². Table 3 shows the difference between these choice proportions. Positive values represent a bias toward the more frequently reinforced option following presentation of the more frequent tone compared to following the less frequent tone.

² Additional analyses for the second half of the trials of the skewed condition revealed unequal proportions of reinforcement, $t(13) = -2.27$, $p < .05$, after the more frequently reinforced option ($M = 2.12$, $SD = .45$) and after the less frequently reinforced option ($M = 2.74$, $SD = .76$), but no correlation with the differences in participants' choice proportions, $r(12) = -.31$, $p = .29$. Note that the unscheduled difference in the proportion of reinforcement was opposite to the difference in participants' choice proportions.

In the skewed condition 10 of 14 participants showed a positive bias. In the criterion-skew and the no-skew conditions positive and negative biases occurred with similar frequencies (6 versus 8 participants, respectively) for both conditions.

A mixed ANOVA revealed a large main effect for distribution type, $F(2, 39) = 57.56$, $p < .001$, and a tone-type-by-distribution-type interaction, $F(2, 39) = 3.27$, $p < .05$. Figure 2 shows the average conditional choice proportions for the more frequently reinforced option in each group.

The main effect reflects participants' sensitivity to the base rate of reinforcement. When the base rate was skewed, participants' mean choice proportion was .83 ($SD = .10$) as compared to a proportion of .51 ($SD = .07$) when reinforcement was evenly distributed (the no-skew condition). There was no statistically significant difference [$t(26) = .53$] between the skewed condition ($M = .85$, $SD = .12$) and the criterion-skewed condition ($M = .82$, $SD = .08$). The significant interaction was due to the fact that, in the skewed condition, the choice proportion for the more frequently reinforced option was higher, $t(13) = 2.26$, $p < .05$, after the more frequent tone ($M = .91$, $SD = .07$) than after the less frequent tone ($M = .78$, $SD = .22$). No such difference was found in the criterion skewed condition, $t(13) = -.60$, $p > .50$; $M = .81/.83$, $SD = .10/.10$, or in the no-skew condition, $t(13) = -.50$, $p > .50$; $M = .50/.52$, $SD = .11/.14$.

These results replicate and extend those of Experiment 1. Again, a contingency between

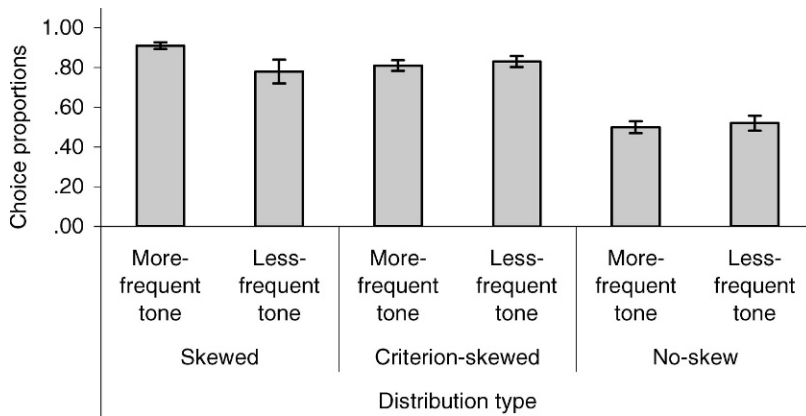


Fig. 2. Mean conditional choice proportions for the more frequently reinforced option as a function of the type of tone (more frequent or less frequent) on the trial. Error bars represent the standard error of the mean.

predictors and criterion events emerged in participants' performance. This pattern is consistent with the phenomenon of a PC because participants related predictors and criterion events only if they were similar in terms of their base rates, that is, occurring more frequently or less frequently.

In a final experiment, we addressed the question of whether associative learning is sufficient to account for the stronger tendency to choose the more frequently reinforced option after the more frequent tone. In other words, we asked whether the greater number of training trials alone can account for the fact that the criterion-event base rate exerts a stronger influence after the more frequent predictor or whether this is more properly a PC because its origin is the similarity of predictor and criterion-event base rates.

Information about the criterion-event base rate is only available conditional on the predictors. This implies less training (e.g., 40 trials compared to 120 trials) for the extraction of the criterion-event base rate presented after the less frequent predictor. In this case, associative learning theories (e.g., Rescorla & Wagner, 1972) predict that after the less frequent predictor, the influence of the skewed criterion-event base rate should be less pronounced than after the more frequent predictor. However, because the learning increment typically decreases with every additional trial, this difference should decrease with extended training. We adopted this logic in Experiment 3 and hypothesized that substantially increasing the amount of training

would not lead to a reduction in the PC because it is supposed to rely on the similarity of the base rates.

EXPERIMENT 3

Experiment 3 was a replication of the skewed condition of Experiments 1 and 2. However, there were 320 trials instead of 160. The less frequent predictor was present on 80 trials, the more frequent on 240 trials.

METHOD

Participants and Design

Twenty-one undergraduate psychology students (14 female, 7 male) from the University of Heidelberg participated. They signed up on a placard promoting an information-processing experiment that offered an ordinary performance-contingent reward.

Apparatus

Up to 6 students participated at the same time in an artificially lit laboratory located within the psychology department. They were seated in cubicles equipped with a personal computer without being able to observe the other participants. The tones were delivered via earphones, and participants could adjust the volume to their liking.

Procedure

Computers delivered the instructions, controlled stimulus presentation, and recorded

Table 3

Bias towards choosing the more frequently reinforced option at a higher rate after the more frequent tone than after the less frequent tone in Experiment 2.

Conditions	Skewed		Criterion-skewed		No-skew	
Individual Data	-.11	(.84, .95)	-.20	(.72, .92)	-.26	(.32, .58)
	-.08	(.90, .98)	-.17	(.69, .87)	-.26	(.50, .76)
	.00	(.90, .90)	-.16	(.72, .88)	-.21	(.39, .60)
	.00	(1.00, 1.00)	-.16	(.82, .98)	-.20	(.32, .52)
	.02	(.92, .91)	-.12	(.61, .73)	-.17	(.50, .67)
	.03	(1.00, .97)	-.06	(.89, .95)	-.16	(.47, .64)
	.05	(.88, .83)	-.02	(.98, .00)	-.07	(.43, .50)
	.06	(1.00, .94)	-.02	(.79, .81)	-.01	(.46, .46)
	.17	(.84, .68)	.04	(.91, .87)	.02	(.54, .52)
	.17	(.84, .66)	.05	(.80, .76)	.03	(.62, .59)
	.23	(.87, .64)	.10	(.84, .74)	.07	(.50, .43)
	.25	(.85, .60)	.11	(.87, .76)	.14	(.65, .51)
	.35	(.97, .63)	.16	(.92, .76)	.30	(.68, .38)
	.79	(.99, .20)	.17	(.85, .68)	.41	(.59, .17)
Mean	.14		-.02		-.03	
SD	.23		.13		.20	

Note. See Table 1 for additional details.

participants’ choices. Because of the increased number of trials, we shortened the time the keyboard was locked after the presentation of the tone to 0.5 s. Participants were instructed to maximize their returns. Only one session was conducted per participant.

For every participant, the computer generated 10 random sequences of 32 tones and corresponding correct choices by drawing without replacement from the predetermined distribution shown in Table 4. Again, one of the tones was three times as frequent as the other tone, and pressing one of the keys was reinforced three times as frequently as the other key.

RESULTS AND DISCUSSION

Means of choice proportions for the more frequently reinforced option conditional on the type of tone were calculated for the second 80 trials, as in the previous experiments, and

Table 4

Contingency tables indicating the predetermined distributions used in Experiment 4.

		Reinforced option		
		Left Key	Right Key	
Predictor	High-pitch tone	18	6	24
	Low-pitch tone	6	2	8
		24	8	32

Note. The base rates of both predictors (tones) and criterion events (reinforced options) were skewed at a ratio of 3:1.

for the last 80 trials of the experiment. Table 5 shows the individual differences between the conditional choice proportions. Positive values represent a bias for a stronger tendency to choose the more frequently reinforced option following presentation of the more frequent tone compared to following the less frequent tone. After the second 80 trials 10 of 21 participants showed a positive bias, and after the last 80 trials 13 of 21 participants showed a positive bias. Across participants, the bias for the second 80 trials correlated with the bias for the last 80 trials, $r(21) = .62, p < .05$, showing a considerable degree of stability.

A repeated measures ANOVA revealed a main effect for type of tone, $F(1,20) = 5.00, p < .05$, no main effect for the location of the trials within the experiment, $F(1,20) < 1$ and no interaction, $F(1,20) < 1$. The main effect for the tone type reflects participants’ tendency to form a PC. Figure 3 shows the average choice proportions for the more frequently reinforced option in the middle and at the end of the experiment.

After the second 80 trials, the conditional choice proportion for the more frequently reinforced option was higher, $t(20) = 1.80, p = .09$, after the more frequent tone ($M = .85, SD = .12$) than after the less frequent tone ($M = .76, SD = .24$). The same was still true after the last 80 trials, $t(20) = 2.40, p < .05$, after the more frequent tone, ($M = .87, SD = .14$) than after the less frequent tone ($M = .79, SD = .18$).

Table 5

Individual means of the biases favoring the more frequently reinforced option after the more frequent tone over the same option after the less frequent tone in Experiment 3.

Trials 81 to 160		Trials 241 to 320	
.05	(.85, .80)	-.22	(.68, .90)
-.05	(.90, .95)	-.05	(.95, 1.00)
-.15	(.85, 1.00)	-.05	(.90, .95)
.00	(.50, .50)	-.03	(.57, .60)
.13	(.72, .59)	-.02	(.65, .67)
-.02	(.98, 1.00)	.00	(1.00, 1.00)
.00	(1.00, 1.00)	.00	(1.00, 1.00)
.22	(.69, .48)	.00	(1.00, 1.00)
-.05	(.76, .81)	.01	(.74, .73)
.01	(.90, .89)	.02	(.92, .90)
-.05	(.95, 1.00)	.03	(.88, .85)
.09	(.93, .84)	.04	(.98, .95)
-.02	(.88, .90)	.10	(.95, .85)
-.13	(.82, .95)	.11	(.95, .84)
.05	(.80, .75)	.13	(.63, .50)
-.16	(.79, .95)	.15	(.95, .80)
-.10	(.80, .90)	.15	(.84, .68)
.75	(1.00, .25)	.20	(.93, .74)
.42	(.88, .46)	.34	(.97, .63)
.49	(.97, .48)	.35	(.80, .45)
.50	(.87, .37)	.48	(.95, .47)

Note. The data are ordered from the participant with lowest bias to the participant with the highest bias over last 80 trials. See Table 1 for additional details.

Experiment 3 demonstrates the stability of the stronger tendency to choose the more frequently reinforced option after the more frequent tone over extended training. Its magnitude after the second 80 trials was virtually the same as at the end of training at 320 trials. Although it remains possible that still more training would have eliminated the

effect, the stability of the absolute level of choices, especially after the less frequent tone ($t(21) = .72, p = .48$, comparing the proportion for the second 80 and last 80 trials), suggests that the associative learning process had reached an asymptote. Therefore, it seems unlikely that the number of training trials alone can account for the results. Rather, participants seem to have demonstrated a PC, relating predictors and criterion events on the basis of similar base rates.

GENERAL DISCUSSION

In an operant analysis of what has been conceived as base-rate neglect, Goodie and Fantino (1996) showed that superficial similarity between statistically unrelated predictors and criterion events produced contingency-based predictions, thereby reducing the weight given to criterion-event base rates. Supplementing these findings, we have demonstrated how base rates themselves can produce contingency-based predictions from predictors that are statistically and superficially unrelated to criterion events.

In the critical condition of the three experiments, where the base rate of the criterion events and the base rate of the predictors were both skewed, the similarity of the base rates prompted contingency-based predictions. Specifically, when presented with a more frequent predictor, participants more strongly preferred the more frequently reinforced response option, representing the

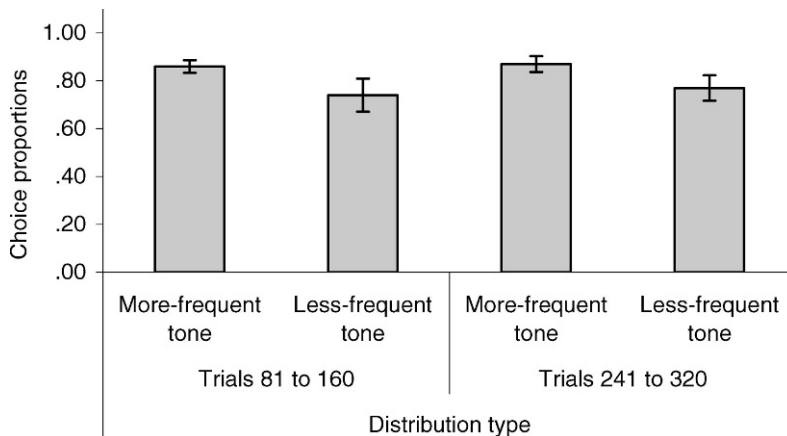


Fig. 3. Mean conditional choice proportions for the more frequently reinforced option as a function of the type of tone (more frequent or less frequent) on the trial. Error bars represent the standard error of the mean.

more frequent criterion event, than when presented with a less frequent predictor.

These results show base-rate neglect in the sense that criterion-event base rates are not used in accordance with Bayesian standards. In the present experiments, performance consistent with those standards would have been only controlled by criterion-event base rates and not by the predictors. Consistent with a tendency to “optimize” predictions (Shanks *et al.*, 2002), the influence of the criterion-event base rates was apparent in the tendency to prefer the option representing the more frequent criterion event. Consistent with a PC-based interpretation (Fiedler & Freytag, 2004; Fiedler *et al.*, 2007), predictor base rates qualify the control by the criterion-event base rates. This relation between predictor and criterion event is a function of the similarity of their base rates. Thus, one set of base rates is responsible for the non-normative impact of the other set of base rates.

There are several implications of these results. On one hand, we add base-rate similarity as a source of behavioral base-rate neglect. On another, the results extend the generality of the PC phenomenon to a behavioral task such as matching-to-sample. Beyond associative learning, there is an additional process triggered by the presence of relatively frequent and infrequent events in the environment. Analogous to the reasoning of Goodie and Fantino (1996), this process could reflect adult humans’ history of predicting on the basis of perceived similarity, here base-rate similarity.

REFERENCES

- Allan, L. G. (1993). Human contingency judgments: Rule based or associative? *Psychological Bulletin*, 114, 435–448.
- Fiedler, K., & Freytag, P. (2004). Pseudocontingencies. *Journal of Personality and Social Psychology*, 87, 453–467.
- Fiedler, K., Freytag, P., & Unkelbach, C. (2007). Pseudocontingencies in a simulated classroom. *Journal of Personality and Social Psychology*, 92, 665–677.
- Goodie, A., & Fantino, E. (1996). Learning to commit or avoid the base-rate error. *Nature*, 380, 247–249.
- Goodie, A., & Fantino, E. (1999). Base rates versus sample accuracy: Competition for control in human matching to sample. *Journal of the Experimental Analysis of Behavior*, 71(2), 155–169.
- Humphreys, L. G. (1939). Acquisition and extinction of verbal expectations in a situation analogous to conditioning. *Journal of Experimental Psychology*, 25, 294–301.
- Rescorla, R. A., & Wagner, A. R. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and non-reinforcement. In A. H. Black, & W. F. Prokasy (Eds.), *Classical conditioning II: Current theory and research* (pp. 64–99). New York: Appleton-Century-Crofts.
- Shanks, D. R. (1990). Connectionism and the learning of probabilistic concepts. *The Quarterly Journal of Experimental Psychology A: Human Experimental Psychology*, 42, 209–237.
- Shanks, D. R., Tunney, R. J., & McCarthy, J. D. (2002). A re-examination of probability matching and rational choice. *Journal of Behavioral Decision Making*, 15, 233–250.
- Tversky, A., & Kahneman, D. (1982). Evidential impact of base-rates. In D. Kahneman, P. Slovic, & A. Tversky (Eds.), *Judgment under uncertainty: Heuristics and biases* (pp. 153–160). Cambridge: Cambridge University Press.

Received: July 10, 2007

Final Acceptance: March 31, 2008